

IN THE CLAIMS

RECEIVED
CENTRAL FAX CENTER
APR 28 2008

Please amend claims 4-6 as follows:

1. (PREVIOUSLY PRESENTED) A method of optimizing a query in a computer system, the query being performed by the computer system to retrieve data from a database stored on the computer system, the method comprising:

(a) during compilation of the query, maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite;

(b) at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns, and generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels; and

(c) performing the query execution plan to retrieve data from a database stored on the computer system.

2. (PREVIOUSLY PRESENTED) The method of claim 1, further comprising:

(1) at query execution time, dynamically determining a grouping sets sequence for the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations based on intermediate grouping sets, in order to optimize the grouping sets sequence.

3. (PREVIOUSLY PRESENTED) The method of claim 2, wherein the dynamically determining step further comprises (1) performing a GROUP BY for a base grouping set and then optimizing execution of the grouping sets sequence by selecting a grouping set having lowest cardinality from a previous one of the levels as an input to a grouping set on a next one of the levels, and (2) performing a UNION ALL operation on the grouping sets.

4. (CURRENTLY AMENDED) A computer-implemented apparatus for optimizing a query, the query being performed to retrieve data from a database, the apparatus comprising:

(a) a computer system;

(b) logic means, performed by the computer system, for

(1) during compilation of the query, maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite;

(2) at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns, and generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels; and

(3) performing the query execution plan to retrieve data from a database stored on the computer system.

5. (CURRENTLY AMENDED) The apparatus of claim 4, further comprising logic means for:

(1) at query execution time, dynamically determining a grouping sets sequence for the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations based on intermediate grouping sets, in order to optimize the grouping sets sequence.

6. (CURRENTLY AMENDED) The apparatus of claim 5, wherein the logic means for dynamically determining step further comprises logic means for (1) performing a GROUP BY for a base grouping set and then optimizing execution of the grouping sets sequence by selecting a grouping set having lowest cardinality from a previous one of the levels as an input to a grouping set on a next one of the levels, and (2) performing a UNION ALL operation on the grouping sets.

7. (PREVIOUSLY PRESENTED) An article of manufacture comprising a program storage device embodying instructions that, when read and executed by a computer system, cause the computer system to perform a method for optimizing a query, the query being performed by the computer system to retrieve data from a database stored in a data storage device coupled to the computer system, the method comprising:

(a) during compilation of the query, maintaining a GROUP BY clause with one or more GROUPING SETS, ROLLUP or CUBE operations in its original form, instead of rewriting the GROUP BY clause, until after query rewrite;

(b) at a later stage of query compilation, translating the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations into a plurality of levels, wherein each of the levels has one or more grouping sets comprised of grouping columns, and generating a query execution plan for the query with a super group block having an array of pointers, wherein each pointer points to the grouping sets for a particular one of the levels; and

(c) performing the query execution plan to retrieve data from a database stored on the computer system.

8. (PREVIOUSLY PRESENTED) The article of manufacture of claim 7, further comprising:

(1) at query execution time, dynamically determining a grouping sets sequence for the GROUP BY clause with the GROUPING SETS, ROLLUP or CUBE operations based on intermediate grouping sets, in order to optimize the grouping sets sequence.

9. (PREVIOUSLY PRESENTED) The article of manufacture of claim 8, wherein the dynamically determining step further comprises (1) performing a GROUP BY for a base grouping set and then optimizing execution of the grouping sets sequence by selecting a grouping set having lowest cardinality from a previous one of the levels as an input to a grouping set on a next one of the levels, and (2) performing a UNION ALL operation on the grouping sets.